ILP Modulo Data

Panagiotis Manolios, Vasilis Papavasileiou, and Mirek Riedewald
Northeastern University
{pete,vpap,mirek}@ccs.neu.edu

arXiv:1404.5665v3 [cs.LO] 15 Sep 2014

***Abstract*—The vast quantity of data generated and captured every day has led to a pressing need for tools and processes to organize, analyze and interrelate this data. Automated reasoning and optimization tools with inherent support for data could enable advancements in a variety of contexts, from data-backed decision making to data-intensive scientific research. To this end, we introduce a decidable logic aimed at database analysis. Our logic extends quantifier-free Linear Integer Arithmetic with operators from Relational Algebra, like selection and cross product. We provide a scalable decision procedure that is based on the BC(T) architecture for ILP Modulo Theories. Our decision procedure makes use of database techniques. We also experimentally evaluate our approach, and discuss potential applications.**

I. Introduction

In 2010, enterprises and users stored more than 13 exabytes of new data [1]. Database Management Systems (DBMS's) based on the Relational Model [3] are a key component in the computing infrastructure of virtually any organization. With big data playing a determining role in business and science, we are motivated to rethink data management and analysis.

Database systems capable of symbolic computation could enable powerful new methodologies for strategic planning, decision making, and scientific research. We propose database systems that (a) store symbolic (in addition to concrete) data, and at the same time (b) allow queries of a symbolic nature, *e.g.*, with free variables. Such database systems can be dually thought of as constraint solvers that reason in the presence of data. Symbolic data allows us to encode partially specified or entirely speculative information, *e.g.*, database entries that exist for the purpose of what-if analysis. Symbolic queries enable deductive reasoning about data.

Existing relational query languages (*e.g.*, SQL) only allow concrete data and queries. Symbolic enhancements require a formalism that combines constraints and relational queries. We address this need by introducing the Δ logic. Δ extends quantifier-free Linear Integer Arithmetic (QFLIA) with database tables and operators from Relational Algebra, like selection (σ), union ($\cup$), and cross product ($\times$). While Δ is decidable, the logic in its general form gives rise to hard satisfiability problems, primarily because it allows universal quantification over cross products of big tables. We study unrestricted Δ (for it is a natural umbrella formalism), but also provide restrictions that enable an efficient decision procedure. In other words, we identify a class of database problems that are a realistic initial target for formal analysis.

This research was supported in part by DARPA under AFRL Cooperative Agreement No. FA8750-10-2-0233 and by NSF grants CCF-1117184 and CCF-1319580.

We provide a scalable procedure based on the BC(T) architecture for ILP Modulo Theories (IMT) [10]. Our approach is dubbed *ILP Modulo Data*, because an ILP solver co-exists with a procedure that establishes a correspondence between integer variables and database tables. The latter contain a mix of concrete and symbolic data. ILP Modulo Data allows us to use a powerful ILP solver based on branch-and-cut (B&C) on the arithmetic side, while also utilizing database techniques that allow us to scale to realistic datasets.

The compositional nature of ILP Modulo Data is well-suited for potential applications. Organizations have access to vast amounts of data, but at the same time rely heavily on Mathematical Programming technology. We enhance Mathematical Programming tools with the ability to directly access data, thus assisting data-backed decision making. Such tools would also benefit scientists in fields ranging from ornithology [16] to astronomy [5], by providing immediate feedback on the consistency between models the scientists devise and datasets of observations they collect. Our paper outlines potential applications, while our experimental evaluation relies on benchmarks that characterize them. We experimentally demonstrate that our ILP Modulo Data framework provides better performance than the approach of eagerly reducing Δ to QFLIA.

Paper Structure: Section II introduces our reasoning paradigm through a motivating example. Section III presents the Δ logic, while Section IV identifies a Δ fragment that yields scalable procedures. Section V describes our decision procedure. We experimentally evaluate our approach in Section VI. We provide an overview of related work in Section VII, and conclude with Section VIII.

II. Motivating Example

Our motivating example (formalized in Figure 1) concerns the problem of optimally investing a given amount of capital. This is an appropriate application for our techniques, because (a) investments are almost always data-driven as they take historical stock prices into account, and (b) financial institutions already rely on Mathematical Programming.

The problem involves investing in a *portfolio* of n publicly traded stocks, with the goal of maximizing profit while following guidelines that minimize risk. A database provides information on these stocks, including stock prices from the New York Stock Exchange (NYSE). We would like to pick the n stocks that would have yielded the highest profit over a period of time in the recent past, *e.g.*, over the preceding year. This optimization problem is subject to risk-mitigation constraints that require us to pick companies from a variety of sectors. While investing in the exact solver-generated portfolio

Id	Cap	Sector
1 (EMC)	large	tech
2 (FII)	medium	financials
3 (AKR)	small	retail
...	...	...

(a) `stocks`

Id	Diff
1	128
2	117
3	89
...	...

(b) `quotes`

maximize
$\Sigma_{1\leq i\leq n} a_i \cdot d_i$
subject to

$(x_i, c_i, s_i) \in \texttt{stocks}$,	$1 \leq i \leq n$
$(x_i, d_i) \in \texttt{quotes}$,	$1 \leq i \leq n$
$x_i \neq x_j$,	$1 \leq i < j \leq n$
$\Sigma_{\{i \mid 1\leq i\leq n, s_i = s\}} a_i \leq \Sigma_{1\leq i\leq n} a_i/3$,	for every sector s
$\Sigma_{\{i \mid 1\leq i\leq n, c_i = \mathsf{small}\}} a_i \leq \Sigma_{1\leq i\leq n} a_i/4$	

(c) Constraints

Fig. 1. Portfolio Management with ILP Modulo Data

(which relies only on past performance) is not necessarily a good strategy, such a portfolio provides useful information for the analysts who make the final investment decisions.

The data is given in tables `stocks` and `quotes` (Figures 1a and 1b). Each company in `stocks` is described by a unique ID (with the associated NYSE symbol parenthesized), its capitalization (small, medium, or large), and its sector (*e.g.*, tech, retail, financials, automotive, energy, emerging-markets). While Figure 1 uses human-readable names, we can encode these fields with bounded integer quantities. Each entry in `quotes` describes the observed movement of a certain stock in a given timeframe, assuming that dividends were reinvested. For example, the first row describes an increase of 28% in the price of EMC. `quotes` is an application-specific abstraction, *i.e.*, the actual database contains past stock prices and `quotes` is a *view* produced by comparing data for two time periods.

The i^{th} stock in the portfolio is characterized by a unique ID x_i that corresponds to entries in the dataset, *i.e.*, there exist entries $(x_i, c_i, s_i) \in \texttt{stocks}$ and $(x_i, d_i) \in \texttt{quotes}$. To minimize risk, we force the n IDs x_i to be distinct, and allow no single sector to account for more than a third of the total capital. Additionally, no more than a fourth of the capital goes to smallcap companies. The objective function maximizes the capital at the end of the period, and thus the profit.

Note that if the amounts a_i are variables, the objective function is non-linear. The problem can be circumvented by providing integer constants for a_i, *i.e.*, by specifying how the capital will be partitioned. With constants for a_i, the non-table constraints are essentially in QFLIA. (The summations for i that satisfy conditions like $s_i = s$ and $c_i = \mathsf{small}$ are easy to encode as sums of if-then-else terms.) Conversely, the problem is essentially satisfiability of an arithmetic instance, where certain variables correspond to database contents. This is the kind of problem that we propose new techniques for. We cannot use a standalone DBMS, since DBMS's do not handle constraints and optimization. Neither are existing solvers up to the task, since they do not provide ways of managing data.

The constraints we have described are meant to be representative. Clearly investors also have to consider other options, including investing in index funds, bonds, debt securities and derivative contracts. These financial instruments may have other characteristics that need to be modeled. Our constraints are also based on simplifying assumptions, *e.g.*, that we can invest an arbitrary amount in any given stock at any time. It is not within the scope of our paper to model investment problems comprehensively. What matters is that these additional concerns also mix arithmetic with data, thus reinforcing the need for data-aware solving.

III. The Logic Δ

$$F ::= T_1 \leq T_2 \mid \exists D \mid \neg F \mid F_1 \vee F_2$$

$$D ::= \{T^+\} \mid \langle \sigma\, x : F : D \rangle \mid D_1 \times D_2 \mid D_1 \cup D_2$$

$$T ::= (T_1, T_2) \mid \mathsf{left}(T) \mid \mathsf{right}(T) \mid x \mid K \mid K \cdot T \mid T_1 + T_2$$

$$K ::= \ldots \mid -2 \mid -1 \mid 0 \mid 1 \mid 2 \mid \ldots$$

Fig. 2. Grammar of Δ

This Section introduces the logic Δ. Δ combines arithmetic with queries over tabular data. Δ thus encompasses database problems like our motivating example of Section II.

The grammar of Δ is given in Figure 2. K, T, D, and F are the non-terminal symbols for integer constants, terms, tables, and formulas, respectively. The first line of productions for T corresponds to pairs and their accessors; the second line is for variable symbols (x) and integer expressions. A table (non-terminal symbol D) is either an *input table*, a *selection*, a *cross-product*, or a *union*. The selection $\langle \sigma\, x : F : D \rangle$ is a table that consists of only those entries in D that satisfy F, *i.e.*, the variable x ranges over the table entries; σ binds x in F, but not in D. For formulas (non-terminal symbol F), $\exists D$ should be read as "D is not empty". All other constructs bear the obvious meaning. We assume that all variables not bound by σ are integer. We will freely use derived operators, *e.g.*, conjunction and integer equality.

Δ is typed. Each term is either of type `int` or of type $s * t$, where s and t are types. left and right are only permissible when applied to a term of type $s * t$ for some type s and some type t; if x is of type $s * t$, then $\mathsf{left}(x)$ is of type s and $\mathsf{right}(x)$ is of type t. The integer constants are of type `int`. The arithmetic operators ($+$, $\cdot$, and $\leq$) only apply to terms of type `int`; $+$ and $\cdot$ produce integers. Each table has a *schema*, which is the type of its entries. (Schemas are the table-level counterpart of types.) An input table is comprised of entries of the same type. If table D_1 has schema s_1 and table D_2 has schema s_2, then $D_1 \times D_2$ has schema $s_1 * s_2$. For $\langle \sigma\, x : F : D \rangle$ to be properly typed, F should be a properly-typed formula under the assumption that the type of x is the schema of D; the schema of $\langle \sigma\, x : F : D \rangle$ is the same as the schema of D. Union expects tables of the same schema and preserves it.

Clearly, Δ is at least as powerful as QFLIA. At the same time, Δ encompasses most features one would expect from a relational query language. We have left out certain operators usually present in query languages. First, note that projection

(π) would not provide additional power, since it is possible to refer to any subset of the columns, without producing an intermediate table that leaves out the irrelevant ones. Also, the set difference $A \setminus B$ can be encoded as $\langle \sigma\, a : \neg\, \exists \langle \sigma\, b : a = b : B \rangle : A \rangle$, assuming that the schema of A and B has exactly one column; otherwise, in place of $a = b$ we would have a conjunction of equalities over all columns. Additionally, Δ can express many forms of aggregation, including count (when compared to a constant), min, and max.

Example 1. *The portfolio encoded by Figure 1 can be represented as the input table*

$$\texttt{portfolio} = \{(1, (x_1, a_1)), \ldots, (n, (x_n, a_n))\}.$$

portfolio *contains symbolic data, something which is not allowed by DBMS's. The first column ensures that the* n *entries are distinct, irrespective of the assignment.* portfolio *is of schema* $\texttt{int} * (\texttt{int} * \texttt{int})$. *Consider the following constraint:*

$$\neg\, \exists \boxed{\begin{array}{lll} \langle \sigma\, x : & \mathsf{left}(\mathsf{left}(x)) \neq \mathsf{left}(\mathsf{right}(x)) \wedge & \\ & \mathsf{left}(\mathsf{right}(\mathsf{left}(x))) = \mathsf{left}(\mathsf{right}(\mathsf{right}(x))) & \\ : & \boxed{\texttt{portfolio}} \times \boxed{\texttt{portfolio}} & \rangle \end{array}}$$

The constraint states that there are no entries $(i, (x_i, a_i))$ *and* $(j, (x_j, a_j))$ *in* portfolio *such that* $i \neq j$ *and* $x_i = x_j$, i.e., portfolio *references* n *distinct stocks (as was our intention in Figure 1). The constraint essentially involves universal quantification over* portfolio $\times$ portfolio.

A. Decidability

Δ satisfiability can be reduced to QFLIA satisfiability. We explain the reduction briefly. We represent a table expression D of schema s as a set $[\![D]\!]$ consisting of pairs $r \oslash b$, where r is a term of type s and b is a QFLIA formula, with the intended meaning that r is present in the table iff b is true. We use the operator $\oslash$ to distinguish the auxiliary pairs used for the reduction from the ones allowed by the syntax of Δ. For a formula F, $[\![F]\!]$ denotes the corresponding formula in QFLIA; similarly for integer terms. $F[x/r]$ stands for substituting x with r in F, with appropriate care for occurrences of the symbol x bound by σ inside F. We define $[\![\cdot]\!]$ for tables and formulas below as two mutually recursive functions.

$$\begin{aligned} [\![\{r_1, \ldots, r_n\}]\!] &= \{r_1 \oslash \mathsf{true}, \ldots, r_n \oslash \mathsf{true}\} \\ [\![\langle \sigma\, x : F : D \rangle]\!] &= \{r \oslash (b \wedge [\![F[x/r]]\!]) \mid r \oslash b \in [\![D]\!]\} \\ [\![D_1 \times D_2]\!] &= \{(r_1, r_2) \oslash (b_1 \wedge b_2) \mid \\ &\qquad r_1 \oslash b_1 \in [\![D_1]\!], r_2 \oslash b_2 \in [\![D_2]\!]\} \\ [\![D_1 \cup D_2]\!] &= [\![D_1]\!] \cup [\![D_2]\!] \end{aligned} \tag{1}$$

$$\begin{aligned} [\![T_1 \leq T_2]\!] &= [\![T_1]\!] \leq [\![T_2]\!] \\ [\![\exists\, D]\!] &= \bigvee_{r \oslash b \in [\![D]\!]} b \\ [\![\neg F]\!] &= \neg [\![F]\!] \\ [\![F_1 \vee F_2]\!] &= [\![F_1]\!] \vee [\![F_2]\!] \end{aligned} \tag{2}$$

For encoding Δ integer terms as QFLIA terms (*e.g.*, $[\![T_i]\!]$ in Equation 2), all that needs to be done is elimination of pair constructors and accessors via the rules $\mathsf{left}((x, y)) = x$ and $\mathsf{right}((x, y)) = y$. The reduction suffices to establish decidability of Δ. The reduction also provides formal semantics for Δ by specifying its meaning in terms of QFLIA.

B. Complexity

Theorem 1. *The satisfiability problem for* Δ *is in NEXPTIME.*

Proof Sketch. The reduction to QFLIA (Equations 1 and 2) produces a formula exponentially larger than the input. Since QFLIA is in NP, the reduction provides a non-deterministic exponential time procedure for Δ-satisfiability. □

Theorem 2. *The satisfiability problem for* Δ *is PSPACE-hard.*

Proof Sketch. We reduce the (PSPACE-complete) QBF problem to Δ satisfiability in polynomial time. We deal with Boolean quantification by quantifying over the input table $\mathcal{B} = \{0, 1\}$. For example, the formula $\forall x \exists y (x \vee \neg y)$ becomes $\neg\, \exists \boxed{\langle \sigma\, x : \neg\, \exists \boxed{\langle \sigma\, y : x = 1 \vee y = 0 : \mathcal{B} \rangle} : \mathcal{B} \rangle}$. □

Complexity analysis of Δ beyond Theorems 1 and 2 is not within the scope of this paper, and has mostly theoretical significance. In practice, query size is orders of magnitude smaller than data size. Conversely, it is meaningful to study *data complexity* [18], *i.e.*, complexity where only the amount of data varies. Instead of assuming a query of constant size, we provide a stronger result by limiting the number of tables that can participate in a cross product. (We also limit nested quantifiers, because the latter can simulate cross products.) We define below the rank function that characterizes this number.

$$\begin{aligned} \mathrm{rank}(\{r_1, \ldots, r_n\}) &= 1 \\ \mathrm{rank}(\langle \sigma\, x : F : D \rangle) &= \mathrm{rank}(F) + \mathrm{rank}(D) \\ \mathrm{rank}(D_1 \times D_2) &= \mathrm{rank}(D_1) + \mathrm{rank}(D_2) \\ \mathrm{rank}(D_1 \cup D_2) &= \max(\mathrm{rank}(D_1), \mathrm{rank}(D_2)) \end{aligned} \tag{3}$$

$$\begin{aligned} \mathrm{rank}(T_1 \leq T_2) &= 0 \\ \mathrm{rank}(\exists\, D) &= \mathrm{rank}(D) \\ \mathrm{rank}(\neg F) &= \mathrm{rank}(F) \\ \mathrm{rank}(F_1 \vee F_2) &= \max(\mathrm{rank}(F_1), \mathrm{rank}(F_2)) \end{aligned} \tag{4}$$

Definition 1 (k-Δ)**.** *For any natural number* k, k-Δ *is the set of formulas* $\{F \mid F \in \Delta$ *and* $\mathrm{rank}(F) \leq k\}$.

Theorem 3. *For any natural number* k, k-Δ *is NP-complete.*

Proof Sketch. k-Δ is NP-hard , because any QFLIA formula can be reduced to a $0-\Delta$ formula in polynomial time ($0-\Delta \subseteq k$-Δ). We obtain membership in NP from the reduction defined by Equation 2, which produces polynomially-sized QFLIA formulas. □

Given the class of formulas k-Δ for some k, the reduction produces QFLIA formulas of size $O(n^{k+1})$, where n is the input size. While the reduction is polynomial (since k is fixed), it may not be practical even for $k = 2$, given that datasets of millions of entries are common. Conversely, we propose restrictions that yield a lazy solving architecture.

IV. The Existential Fragment of Δ

We proceed to study the *existential fragment* of Δ, which we denote by $\exists\Delta$.

Definition 2 ($\exists\Delta$)**.** *A Δ formula belongs to $\exists\Delta$ if the $\exists$ operator always appears below an even number of negations,* i.e.*, $\exists$ only appears with positive polarity.*

The motivation for studying $\exists\Delta$ is as follows. Universal quantification pushes for an approach similar to quantifier instantiation, *e.g.*, Example 1 (which is not in $\exists\Delta$) inherently requires instantiating a constraint for every element in $\texttt{portfolio} \times \texttt{portfolio}$. This can be done incrementally by applying patterns that are standard in verification tools. In contrast, we are not aware of techniques that would be a good match for the kind of existential quantification that arises in Δ. Therefore, the rest of this paper focuses on $\exists\Delta$.

Formulas in $\exists\Delta$ can be transformed into formulas in a convenient intermediate logic without cross products, selections, or unions. We rephrase $\exists$ in terms of a new membership operator. Each formula of the form $\exists\, D$ is viewed as $x \in D$, where $\in$ has the obvious semantics and x is a properly shaped row comprised of fresh integer variables. We will refer to rows like x that serve as witnesses for $\exists$ as *witness rows*. The next step is to translate membership in arbitrary table expressions to membership in input tables. $(x, y) \in D \times E$ becomes $x \in D \wedge y \in E$, while $x \in D \cup E$ becomes $x \in D \vee x \in E$. Finally, $x \in \langle \sigma\, y : F : D \rangle$ becomes $F[y/x] \wedge x \in D$. We eliminate all cross products, selections, and unions by repeated application of the above transformations.

Example 2. *The tables of Figures 1a and 1b can be easily encoded as Δ input tables of schemas* $\texttt{int} * (\texttt{int} * \texttt{int})$ *and* $\texttt{int} * \texttt{int}$*. Let* small *capitalization be represented by the constant* 0*. Consider the following constraint:*

$$\exists \left| \begin{array}{ll} \langle \sigma\, x : & \mathsf{left}(\mathsf{left}(x)) = \mathsf{left}(\mathsf{right}(x)) \wedge \\ & \mathsf{left}(\mathsf{right}(\mathsf{left}(x))) = 0 \wedge \\ & \mathsf{right}(\mathsf{right}(x)) \geq 150 \\ : & \boxed{\texttt{stocks}} \times \boxed{\texttt{quotes}} \qquad \rangle \end{array} \right.$$

The constraint asserts the existence of some tuple $((x_1, (x_2, x_3)), (x_4, x_5)) \in \texttt{stocks} \times \texttt{quotes}$ that satisfies $\Phi = [x_1 = x_4 \wedge x_2 = 0 \wedge x_5 \geq 150]$. (We have eliminated the accessors left *and* right*.) This is equivalent to asserting that $(x_1, (x_2, x_3)) \in \texttt{stocks} \wedge (x_4, x_5) \in \texttt{quotes} \wedge \Phi$.*

The procedure we outlined produces a *decomposed* formula consisting of a QFLIA part and *membership constraints*. We proceed to define these notions formally.

Definition 3 ((Conditional) Membership Constraint)**.** *A membership constraint is a constraint of the form*

$$(x_1, \ldots, x_k) \in \{(y_{1,1}, \ldots, y_{1,k}), \ldots, (y_{l,1}, \ldots, y_{l,k})\} \quad (5)$$

for positive integers k and l and variable symbols x_i, $y_{j,i}$. A constraint of the form $b = 1 \Rightarrow m$, where b is a variable symbol and m is a membership constraint, is called a conditional *membership constraint.*

A membership constraint may hold conditionally, either because it arises from an $\exists$-atom that appears under propositional structure (and therefore holds conditionally), or because of a disjunction introduced by the union operator. We use conditions of the form $b = 1$ because ILP necessitates $[0, 1]$-bounded integer variables in place of Boolean variables. Implication in the opposite direction is never needed, since $\exists$ always appears with positive polarity (as per Definition 2).

Membership constraints do not contain arbitrary arithmetic expressions, but only variable symbols. "Variable abstraction" [9] eliminates richer expressions. While variable abstraction allows for compositional reasoning and helps with theoretical analysis, a limited fragment of arithmetic in membership constraints yields more efficient implementation. Part of our discussion will involve tables that contain integer constants and terms of the form $v + c$, where v is a variable symbol and c is an integer constant. (Everything we present is easy to generalize for such terms.) For convenience, we flatten out rows constructed using the pair constructor of Figure 2, and instead deal with k-tuples of integers. This is only a matter of presentation and has no impact on the algorithms.

Definition 4. *A* decomposed *formula is a conjunction $F \wedge M$, where (a) F is a QFLIA formula and (b) M is a conjunction of possibly conditional membership constraints.*

Theorem 4. *$\exists\Delta$ satisfiability is NP-complete.*

Proof. $\exists\Delta$ satisfiability is NP-hard, because $\exists\Delta$ is at least as powerful as QFLIA. $\exists\Delta$ satisfiability is in NP, because we can reduce $\exists\Delta$ to QFLIA in polynomial time. The reduction first produces a formula in decomposed form (Definition 4). Equation 5 is equivalent to $\bigvee_{j=1,\ldots,l} \bigwedge_{i=1,\ldots,k} x_i = y_{j,i}$; therefore, the membership operator can be eliminated. The result is a formula in QFLIA. □

The polynomial size of the reduction relies on the fact that Δ does not allow tables to be named and referenced from multiple places, *i.e.*, table expressions are not DAG-shaped. Despite the polynomial reduction, a lazy scheme remains relevant. The reason is that QFLIA solvers are not meant for long disjunctions that essentially encode database tables.

V. BC(T) for Δ

The decomposed form of Definition 4 is particularly suited for a scheme that combines separate procedures for QFLIA and table membership. Given that the QFLIA part can be encoded as a conjunction of integer linear constraints [10], it becomes possible to solve instances in decomposed form (and by extension $\exists\Delta$ instances) by instantiating the BC(T) framework for IMT [10]. An ILP solver deals with the QFLIA constraints, and exchanges information with a procedure that checks membership in finite sets. Since database queries typically have simple propositional structure, we do not expect encoding the latter with linear constraints to be a bottleneck.

The membership procedure is confronted with a conjunction of membership constraints (Definition 3). Dealing with conditional constraints is essentially a matter of Boolean search. The membership procedure needs to understand equality atoms,

equality being a primitive. (Our setting is standard first-order logic with equality.) In particular, the procedure keeps track of truth assignments to the equalities in:

$$\{x_i = y_{j,i} \mid j \in [1,l], i \in [1,k]\} \tag{6}$$

The symbols x_i and $y_{j,i}$ have the same meaning as in Definition 3. In the presence of multiple membership constraints, the union of sets, like in Equation 6, is relevant. Given that membership constraints can be checked in isolation, our discussion proceeds with a single constraint. The variables x_i and $y_{j,i}$ also appear in linear constraints. It simplifies our design to assume that all of them appear in ILP, even if they are unconstrained there. The BC(T) framework provides a mechanism ("difference constraints" [10]) for notifying background procedures about atoms like the ones in Equation 6. Given truth values for these atoms, we check that a membership constraint is satisfied by simply traversing the table and looking for a tuple that is column-wise equal to the witness row. The constraint is violated if for every $j \in [1,l]$, there exists some $i \in [1,k]$ such that $x_i \neq y_{j,i}$, *i.e.*, there is no candidate tuple.

The arithmetic and membership parts share variables. It is vital that we systematically explore the space of (dis)equalities between these variables. This exchange of information resembles the non-deterministic Nelson-Oppen scheme (NO) for combining decision procedures [14]. We demonstrate that NO can accommodate membership constraints.

Definition 5 (Arrangement)**.** *Let E be an equivalence relation over a set of variables V. The set*

$$\alpha(V,E) = \{x = y \mid xEy\} \cup \{x \neq y \mid x,y \in V \text{ and not } xEy\}$$

is the arrangement *of V induced by E.*

Definition 6 (Stably-Infinite Theory)**.** *A Σ-theory T is called stably-infinite if for every T-satisfiable quantifier-free Σ-formula F there exists an interpretation satisfying $F \wedge T$ whose domain is infinite.*

Fact 1 (Nelson-Oppen for Stably-Infinite Theories [14, 9])**.** *Let T_i be a stably-infinite Σ_i-theory, for $i = 1,2$, and let $\Sigma_1 \cap \Sigma_2 = \emptyset$. Also, let Γ_i be a conjunction of Σ_i-literals. $\Gamma_1 \cup \Gamma_2$ is $(T_1 \cup T_2)$-satisfiable iff there exists an equivalence relation E of the variables V shared by Γ_1 and Γ_2 such that $\Gamma_i \cup \alpha(V,E)$ is T_i-satisfiable, for $i = 1,2$.*

Lemma 1 (Nelson-Oppen with Propositional Structure)**.** *Let T_i be a stably-infinite Σ_i-theory, for $i = 1,2$, and let $\Sigma_1 \cap \Sigma_2 = \emptyset$. Also, let Φ_i be a quantifier-free Σ_i-formula. $\Phi_1 \wedge \Phi_2$ is $(T_1 \cup T_2)$-satisfiable iff there exists an equivalence relation E of the variables V shared by Φ_1 and Φ_2 such that $\{\Phi_i\} \cup \alpha(V,E)$ is T_i-satisfiable, for $i = 1,2$.*

Proof.

- If $\Phi_1 \wedge \Phi_2$ is $(T_1 \cup T_2)$-satisfiable, there exists a first-order model L that satisfies it. The way L interprets the variables in V gives rise to an equivalence relation E over V such that L satisfies $T_i \cup \{\Phi_i\} \cup \alpha(V,E)$, $i = 1,2$.
- If there exists an equivalence relation E over V such that $\{\Phi_i\} \cup \alpha(V,E)$ is T_i-satisfiable, then there exists a model L_i that T_i-satisfies Φ_i, $i = 1,2$. Let

$$\begin{aligned}\Gamma_i = \quad & \{t \mid t \text{ is an atom in } \Phi_i, L_i \models t\} \ \cup \\ & \{\neg t \mid t \text{ is an atom in } \Phi_i, L_i \models \neg t\}, \quad i = 1,2.\end{aligned}$$

$\Gamma_i \cup \alpha(V,E)$ is T_i-satisfiable, $i = 1,2$. By Fact 1, $\Gamma_1 \cup \Gamma_2$ is $(T_1 \cup T_2)$-satisfiable. But $\Gamma_1 \cup \Gamma_2 \models \Phi_1 \wedge \Phi_2$. Therefore, $\Phi_1 \wedge \Phi_2$ is $(T_1 \cup T_2)$-satisfiable.

□

Lemma 2 (Nelson-Oppen for Membership Constraints)**.** *Let T be a stably-infinite Σ-theory. Also, let Γ be a conjunction of Σ-literals, and M be a conjunction of possibly negated membership constraints. $\Gamma \cup M$ is T-satisfiable iff there exists an equivalence relation E of the variables V shared by Γ and M such that $\Gamma \cup \alpha(V,E)$ is T-satisfiable and $M \cup \alpha(V,E)$ is satisfiable.*

Proof. Membership constraints can be viewed as disjunctions of conjunctions (Proof of Theorem 4) in which no function, predicate, and constant symbols appear, *i.e.*, in the empty signature. The theory pertaining the membership constraints is the empty theory ($\emptyset$), since no axioms are needed. $\emptyset$ is trivially stably-infinite. Our proof obligation follows by applying Lemma 1 with $T_1 = T$ and $T_2 = \emptyset$. □

Note that Lemma 2 allows negated membership constraints. While the latter do not pose algorithmic difficulties, our discussion is limited to the positive occurrences needed for $\exists\Delta$. The statement of Lemma 2 is structurally similar to that of Fact 1, with membership constraints replacing the constraints of some participating stably-infinite theory. It follows that a membership procedure can participate in NO as a black box, much like a theory solver, even though we have not formalized membership constraints by means of a theory. We can thus combine a form of set reasoning with any stably-infinite theory.

BC(T) guarantees completeness for the combination of ILP with a stably-infinite theory [10] by ensuring that the branching strategy explores all possible arrangements. We established that membership can be used much like a stably-infinite theory. All that is needed for completeness is a membership procedure capable of checking consistency of its constraints conjoined with a given arrangement (that contains all literals of Equation 6). As we have seen, this operation is simple and involves no arithmetic. In pursuit of efficiency, we proceed to describe branching and propagation techniques based on table contents. Meaningful branching and propagation involve the integer bounds of variables, *i.e.*, necessitate limited arithmetic reasoning on the membership side.

A. Propagation

B&C-based ILP solvers keep track of variable lower and upper bounds, and heavily rely on bounds propagation algorithms. We describe how to enhance such propagation to exploit the structure of membership constraints.

We denote by $\mathrm{lb}(v)$ an $\mathrm{ub}(v)$ the current lower and upper bounds on variable v. $\mathrm{lb}(v)$ (respectively $\mathrm{ub}(v)$) is either an integer constant, or $-\infty$ (resp. $+\infty$) if no bound is known.

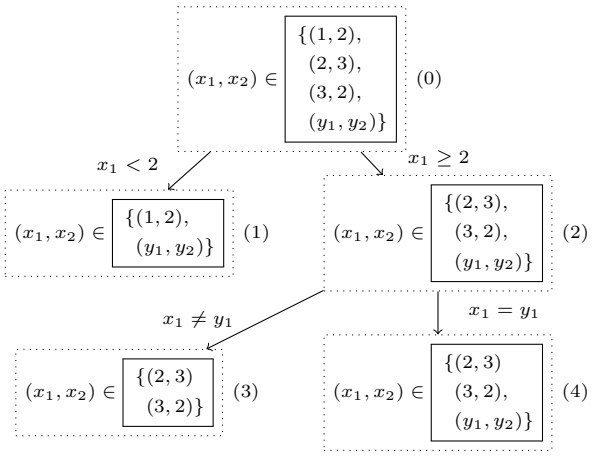

Fig. 3. Data-Driven Branching

We use the notation $\mathrm{lb}'(v)$ and $\mathrm{ub}'(v)$ for bounds on v that the membership procedure infers. We proceed with a membership constraint as per Definition 3. Let $x = (x_1, \ldots, x_k)$; similarly, we denote by y_j the tuple $(y_{j,1}, \ldots, y_{j,k})$. Let $\mathrm{match}(x, y_j)$ be true if and only if for all $i \in [1, k]$, the sets $[\mathrm{lb}(x_i), \mathrm{ub}(x_i)]$ and $[\mathrm{lb}(y_{j,i}), \mathrm{ub}(y_{j,i})]$ intersect.

$$\mathrm{lb}'(x_i) = \max(\mathrm{lb}(x_i), \min\{\mathrm{lb}(y_{j,i}) \mid j \in [1, l], \mathrm{match}(x, y_j)\}) \quad (7)$$

$$\mathrm{ub}'(x_i) = \min(\mathrm{ub}(x_i), \max\{\mathrm{ub}(y_{j,i}) \mid j \in [1, l], \mathrm{match}(x, y_j)\}) \quad (8)$$

We over-approximate the values of the variables x_i by considering all candidate entries (inner $\min$ and $\max$). The outer $\max$ and $\min$ guarantee that we do not weaken bounds. If there exists exactly one value j such that $\mathrm{match}(x, y_j)$, it is sound to deduce the equalities $x_i = y_{j,i}$, for all $i \in [1, k]$. If there is no candidate entry, inconsistency is reported.

Example 3 (Interleaved Propagation)**.** *Consider the decomposed formula* $x = y \wedge (x, y) \in \{(1, 2), (2, 4), (3, 6), (4, 8)\}$. *The formula corresponds to a query over concrete tuples that any DBMS can evaluate in linear time. It is thus vital that our techniques yield acceptable performance. Equations 7 and 8 bound* x *to* $[\min\{1, 2, 3, 4\}, \max\{1, 2, 3, 4\}] = [1, 4]$ *and* y *to* $[\min\{2, 4, 6, 8\}, \max\{2, 4, 6, 8\}] = [2, 8]$. *Given the equality* $x = y$*, ILP propagation deduces that* $x, y \in [2, 4]$*, since* $[2, 4]$ *is the intersection of permissible ranges for* x *and* y*. The membership procedure detects that* match *now only holds for* $(2, 4)$*, and fixes* x *to* 2 *and* y *to* 4*. The ILP solver in turn deduces unsatisfiability, since* $x = y$ *is violated. No branching was needed. Encoding the formula in QFLIA would hide its structure, leading to search. The example generalizes to other lengths and bounded symbolic data.*

B. Branching and Arrangement Search

It follows from Lemma 2 that a branching strategy which exhaustively explores all possible arrangements of the shared variables guarantees completeness. To achieve better performance, we have to branch with the tabular structure of databases in mind, without overlooking symbolic data.

Figure 3 provides an example. The root node (Node 0) describes a single membership constraint, which we assume to be part of a larger decomposed formula. We maintain integer constants in the table, instead of performing variable abstraction which would introduce auxiliary variables for them. According to Equation 6, the membership procedure needs truth assignments for the equalities in $\{x_1 = 1, x_1 = 2, x_1 = 3, x_1 = y_1, x_2 = 2, x_2 = 3, x_2 = y_2\}$. It would not be wise for the search strategy to overlook that this set originates from a table containing numbers, and treat the set members as if they were atomic propositions unrelated to each other.

In our example, branching on the condition $x_1 < 2$ produces two subproblems. Node 1 shows only the tuples that still apply under the condition $x_1 < 2$, *i.e.*, the ones that still satisfy the predicate match; similarly for Node 2. $x_1 < 2$ is a choice informed by the tabular structure. Since 2 as the value of the first column is close to the "middle" of the table, branching on $x_1 < 2$ rules out approximately half of the candidates. (y_1, y_2) is present in both subproblems (Nodes 1 and 2). Branching based on constant bounds is therefore not enough, for we will possibly have to deal with symbolic tuples. Figure 3 demonstrates further branching on $x_1 = y_1$ to determine whether (y_1, y_2) is a suitable witness for the membership constraint.

The example demonstrates the dual nature of the search strategy needed. The problem naturally pushes towards branch-and-bound (which is a restriction of B&C), *e.g.*, branching on $x_1 < 2$ is meaningful. It remains necessary to also branch on equalities between shared variables (*e.g.*, $x_1 = y_1$), just like in any practical implementation of NO. (To be precise, in ILP we would have two separate nodes for $x_1 < y_1$ and $x_1 > y_1$ in place of $x_1 \neq y_1$.) Implementing NO with B&C enables both kinds of branching.

Branching is organically tied to propagation. Initially (Node 0), assuming no previously known bounds for x_1, the table contents only allow us to bound x_1 to the range $[\min(\mathrm{lb}(y_1), 1), \max(\mathrm{ub}(y_1), 3)]$; if y_1 is unbounded, x_1 remains unbounded. The decisions $x_1 \geq 2$ and $x_1 \neq y_1$ (*i.e.*, Node 3) tighten x_1 to $[2, 3]$. We also obtain the range $[2, 3]$ for x_2, *i.e.*, branching on some column potentially leads to propagation across other columns.

C. Discussion

The analysis of this Section indicates that Δ formulas can be decomposed in such a way that a procedure for table lookup assumes part of the workload. BC(T) is particularly suited for implementing such a combination. BC(T) can easily accommodate data-aware propagation (Section V-A) and branching (Section V-B). Our techniques would be harder to implement within a DPLL(T)-style solver [15], given that the toplevel search of DPLL(T) is over the Booleans (and not the integers). A DPLL(T)-based implementation of our techniques would essentially require integrating branch-and-bound in DPLL(T), which is beyond the scope of our work.

The table lookup procedure can be thought of as a small database engine within the solver. The employed database engine can be an actual DBMS, storing the concrete part of tables and possibly bounds on symbolic fields. A DBMS would provide multiple opportunities for improvements. Equations 7 and 8 essentially describe database aggregation, and thus

provide a starting point for the kinds of queries that apply. DBMS queries can be over multiple tables at a time, and can involve conditions other than bounds. As a matter of fact, the match predicate of Equations 7 and 8 can be strengthened with any condition on the data that follows from the formula (*e.g.*, $x = y$ in Example 3), thus computing tighter bounds. Different kinds of database optimizations apply, *e.g.*, materializing queries for better incremental behavior and smarter indexing based on user input.

$\exists\Delta$ (and its decomposed form) formally characterizes a relevant class of problems that can be solved by a compositional scheme which employs a database engine. Our scheme may actually apply to a superset of $\exists\Delta$.

VI. Applications and Experiments

We have implemented support for databases on top of the Inez constraint solver.[1] Inez is our implementation of the BC(T) architecture for IMT on top of the SCIP (M)ILP solver [2]. We refer to the version of Inez that provides database extensions as Inez*DB*. Inez*DB* supports existential database constraints by means of the BC(T)-based combination described in Section V, but also universal quantification by eager instantiation. Inez*DB* (like Inez) additionally supports objective functions.

We have produced a collection of Inez*DB* input files that have the structure we expect in applications. Our benchmark suite is publicly available and can be used as a starting point towards a richer benchmark suite of problems that involve data and constraints.[2] We provide a brief overview of the application areas that inspire our benchmarks.

A. How-To Analysis

Research in the general direction of reverse data management [11] proposes ways of obtaining the desired results out of a database query. We outline this class of problems through an example, which gives rise to some of our benchmarks.

Example 4 (`emp_join.ml`)**.** *The management of a company is surprised to find out that (according to the corporate database) there is no employee younger than 30 whose yearly income exceeds $60000.* Why not *is not obvious, since income is a complicated function of multiple quantities including a base salary, benefits based on age, employee level (junior, middle, or senior), and bonuses.*

The management consults the database administrator on how to *[12] ameliorate the seeming injustice. Together, they explore bonuses that would allow young employees to exceed the $60000 limit. This amounts to synthesizing tuples for the table of bonuses. An alternative is to adjust various parameters in the income computation,* i.e.*, to modify the query instead of the data [17]. This can be done by replacing constants with variables, and letting the solver come up with suitable values.*

B. Test-Case Generation

Test case generation is relevant for databases [19]. A family of benchmarks in our collection demonstrate test data generation by concretizing tables initially containing symbolic data.

Example 5 (`emp_keys.ml`)**.** *The problem involves two tables, named* `incomes` *and* `employees`*.* `incomes` *has an ID column constrained to reference existing entries in* `employees`*,* i.e.*, there is a foreign key constraint.* `incomes` *contains thousands of tuples with symbolic IDs. A satisfying assignment corresponds to a generated database that meets the foreign key constraint, thus serving as meaningful test input.*

C. Scientific Applications

Studying big datasets is a key aspect of scientific research in fields ranging from ornithology [16] to astronomy [5]. To demonstrate the applicability of our techniques, we provide benchmarks inspired by queries that ornithologists perform.

Example 6 (`birds_box.ml`)**.** *An ornithologist wants to see a rare species in person, but has not decided on a good location. She has access to a database of observations. Each observation describes a bird and the geographic coordinates where it was seen. An area can be described as a symbolic rectangle* $B = [\texttt{longitude}_{\texttt{min}}, \texttt{longitude}_{\texttt{max}}] \times [\texttt{latitude}_{\texttt{min}}, \texttt{latitude}_{\texttt{max}}]$*. Our techniques allow the ornithologist to simply ask for* n *observations of the species of interest that lie in* B*. The query effectively concretizes* B*.*

D. Portfolio Management

We experimented with the portfolio optimization example of Section II. Our exact instance (`portfolio.ml`) encodes a more complex variant of the formalization in Section II. An additional table contains stock dividends; dividends are taken into account in the objective function. We tried a range of parameters with a timeout of one hour, and obtained a range of solutions. Notably, picking an optimal portfolio of 5 out of 50 stocks took 161 seconds; 5 out of 4000 stocks took 1510 seconds; and 6 out of 2000 stocks took 1172 seconds. Such table sizes are realistic, given that NYSE lists approximately 2800 companies.

E. Overview of Results

We compare Inez*DB* against an Inez frontend that solves Δ formulas by eagerly translating them to QFLIA via the encoding of Theorem 4. Inez in turn solves QFLIA formulas by reducing them to constraints that SCIP understands. (These constraints are not strictly ILP, since we utilize specialized constraint handlers [2].) We refer to this configuration simply as Inez, since the only addition to Inez is a new frontend. We also produce SMT-LIB versions of our QFLIA formulas, and run them against the latest available version of Z3 (4.3.1).

We provide 8 benchmark generators that allow different modes of operation (*e.g.*, some of them are able to produce both satisfiable and unsatisfiable benchmarks), and are able to output benchmarks with different table sizes. Our input table sizes range from 60 tuples to 640000 tuples. In total, our parameters give rise to 166 benchmarks. We run all three solvers with a timeout of 1800 seconds and a memory limit of 12GB on a machine that provides 2 Intel Xeon X5677 CPUs of 4 cores each and 96GB of RAM. Figure 4 visualizes our experiments. Inez solves 25 satisfiable and 47 unsatisfiable benchmarks. Inez*DB* solves 74 satisfiable and

[1] https://github.com/vasilisp/inez

[2] http://www.ccs.neu.edu/home/vpap/fmcad-2014.html

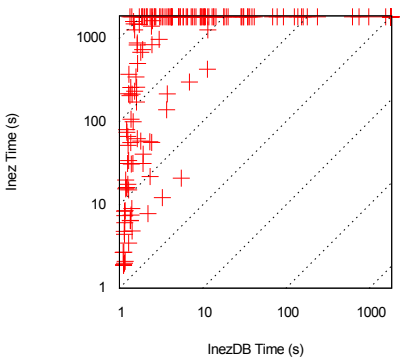

(a) Inez*DB* versus Inez

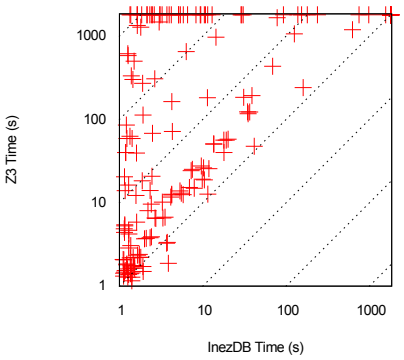

(b) Inez*DB* versus Z3

Fig. 4. Experiments: Inez*DB* versus the eager approach

81 unsatisfiable benchmarks. Finally, Z3 solves 57 satisfiable and 58 unsatisfiable benchmarks. Among the failures for Inez (resp. Z3), 37 (resp. 27) are due to the memory limit. Inez*DB* runs out of memory only once. If we turn off the memory limits, the total numbers of failures don't change much.

Figure 4a indicates that Inez*DB* outperforms Inez by a significant margin. This margin can be attributed to two factors. First, Inez*DB* exploits the structure of database problems (*e.g.*, for branching and propagation), while Inez has no knowledge of this structure. Second, our reduction to QFLIA (in the case of Inez) produces patterns that SCIP is not optimized for, since the latter is designed for MILP and not for QFLIA.

Figure 4b compares Inez against a leading solver for QFLIA (Z3), and thus characterizes the tool's performance in absolute terms. There is a cluster of 40 benchmarks for which Inez*DB* is 2-8 times faster than Z3. (Note that the scale is logarithmic.) Inez*DB* is at least 8 times faster for 31 of the benchmarks that both tools solve, and solves many benchmarks for which Z3 times out. All failures for Inez*DB* are failures for Z3. Z3 outperforms Inez*DB* for only 7 out of the 166 benchmarks, none of which take Inez*DB* more than 4 seconds to solve.

We conclude the evaluation by pointing out that there is significant room for improvement in Inez*DB*. As is the case with almost every first implementation of a new decision procedure, there is room for improvement, *e.g.*, Inez*DB* can benefit from better preprocessing and more sophisticated branching. Inez*DB* can also be improved by adopting database techniques (as we outlined in Section V), or by integrating a DBMS. Our promising experimental results even without such optimizations constitute sufficient evidence that ILP Modulo Data is a viable design for data-enabled reasoning tools.

VII. Related Work

The Constraint Database framework [6] provides a database perspective on constraint solving. The framework encompasses relations described by means of constraints, but not relations comprised of concrete tuples.

"Table constraints" [8, 4], as studied in Constraint Programming, resemble our membership constraints. Such tables are not meant as database tables. Our work differs in significant ways, *e.g.*, our setup allows symbolic table contents. Also, the algorithms presented for table constraints rely on table contents from small domains (*i.e.*, not the reals or the integers). This aligns with the overall emphasis of Constraint Programming, but conflicts with our intended applications.

Veanes et al. describe the Qex technique and tool that uses Z3 to generate tests for SQL queries [19]. Qex essentially encodes the relational operators via axioms, which are later instantiated via E-matching [13]. E-matching is a generic scheme that is not optimized in any way for database problems. Qex is geared towards relatively small tables that suffice as test cases, while our target applications involve bigger tables.

Other approaches tackle constraints arising in database applications with off-the-shelf generic solvers (via eager reductions). Notably, Khalek et al. use Alloy [7], while Meliou and Suciou use MILP [12]. In neither of these approaches does the core of the solver exploit the structure of database instances, *e.g.*, for branching or propagation.

VIII. Conclusions and Future Work

We introduced the ILP Modulo Data framework for marrying data with symbolic reasoning. To that end, we introduced the decidable logic Δ. We identified a fragment of Δ that can be solved efficiently by instantiating the BC(T) architecture. We developed a solver for Δ, and evaluated this solver on a set of benchmarks that we made publicly available.

There are many interesting research directions to be explored in future work, including: (a) the design and implementation of solvers that include an actual DBMS, (b) efficiently handling universal quantification over big tables, say by partitioning input tables and using parallelization, (c) extending our techniques to allow mixed integer, real arithmetic, and other first-order theories, and (d) solving interesting business and scientific applications using the ILP Modulo Data framework.

References

[1] Challenges and Opportunities with Big Data, 2012. Computing Community Consortium White Paper.

[2] Tobias Achterberg. *Constraint Integer Programming*. PhD thesis, Technische Universitat Berlin, 2007.

[3] Edgar Codd. A Relational Model of Data for Large Shared Data Banks. *CACM*, 13(6):377–387, 1970.

[4] Ian Gent, Chris Jefferson, Ian Miguel, and Peter Nightingale. Data Structures for Generalised Arc Consistency for Extensional Constraints. In *AAAI*, 2007.

[5] Jim Gray, Alex Szalay, Ani Thakar, Peter Kunszt, Christopher Stoughton, Don Slutz, and Jan vandenBerg. Data Mining the SDSS SkyServer Database. *arXiv preprint cs/0202014*, 2002.

[6] Paris Kanellakis, Gabriel Kuper, and Peter Revesz. Constraint Query Languages (Preliminary Report). In *PODS*, 1990.

[7] Shadi Abdul Khalek, Bassem Elkarablieh, Yai Laleye, and Sarfraz Khurshid. Query-Aware Test Generation Using a Relational Constraint Solver. In *ASE*, 2008.

[8] Christophe Lecoutre and Radoslaw Szymanek. Generalized Arc Consistency for Positive Table Constraints. In *CP*, 2006.

[9] Zohar Manna and Calogero Zarba. Combining Decision Procedures. In *10th Anniversary Colloquium of UNU/IIST*, 2002.

[10] Panagiotis Manolios and Vasilis Papavasileiou. ILP Modulo Theories. In *CAV*, 2013.

[11] Alexandra Meliou, Wolfgang Gatterbauer, and Dan Suciu. Reverse Data Management. In *VLDB*, 2011.

[12] Alexandra Meliou and Dan Suciu. Tiresias: The Database Oracle for How-To Queries. In *SIGMOD*, 2012.

[13] Leonardo De Moura and Nikolaj Bjorner. Efficient E-matching for SMT solvers. In *CADE-21*, 2007.

[14] Greg Nelson and Derek C. Oppen. Simplification by Cooperating Decision Procedures. *TOPLAS*, 1:245–257, 1979.

[15] Robert Nieuwenhuis, Albert Oliveras, and Cesare Tinelli. Solving SAT and SAT Modulo Theories: From an abstract Davis–Putnam–Logemann–Loveland procedure to DPLL(T). *JACM*, 53(6):937–977, 2006.

[16] Daria Sorokina, Rich Caruana, Mirek Riedewald, Wesley Hochachka, and Steve Kelling. Detecting and Interpreting Variable Interactions in Observational Ornithology Data. In *DDDM*, pages 64–69. IEEE, 2009.

[17] Quoc Trung Tran and Chee-Yong Chan. How to ConQueR Why-Not Questions. In *SIGMOD*, 2010.

[18] Moshe Vardi. The Complexity of Relational Query Languages. In *STOC*, 1982.

[19] Margus Veanes, Nikolai Tillmann, and Peli de Halleux. Qex: Symbolic SQL Query Explorer. In *LPAR-16*, 2010.